

Efficient Creation of Overset Grid Hole Boundaries and Effects of Their Locations on Aerodynamic Loads

William M. Chan*, Shishir A. Pandya†, Stuart E. Rogers‡

NASA Ames Research Center, M/S 258-2, Moffett Field, CA 94035

Recent developments on the automation of the X-rays approach to hole-cutting in overset grids is further improved. A fast method to compute an auxiliary wall-distance function used in providing a first estimate of the hole boundary location is introduced. Subsequent iterations lead to automatically-created hole boundaries with a spatially-variable offset from the minimum hole. An averaged cell attribute measure over all fringe points is used to quantify the compatibility between the fringe points and their respective donor cells for different hole boundary locations. The sensitivity of aerodynamic loads to different hole boundary locations and cell attribute compatibilities is investigated using four test cases: an isolated re-entry capsule, a two-rocket configuration, the AIAA 4th Drag Prediction Workshop Common Research Model (CRM), and the D8 “Double Bubble” subsonic aircraft. When best practices in hole boundary treatment are followed, only small variations in integrated loads and convergence rates are observed for different hole boundary locations.

I. Introduction

IN recent years, structured overset grid technology has been used to perform modeling and simulation analysis on a wide variety of complex aerospace applications.¹⁻⁶ There are three steps in creating a structured overset grid system: surface grid generation, volume grid generation, and domain connectivity. Surface grid generation involves the decomposition of the surface geometry into quadrilateral overlapping blocks, and the distribution of grid points to properly resolve the geometry and flow features. Volume grid generation typically consists of creating near-body curvilinear grids around the surface geometry, and off-body Cartesian grids that fill the region away from the surface. Since both the surface and volume grids are allowed to overlap arbitrarily, grid points that fall inside solid bodies or outside the computational domain need to be identified and excluded from the process of solving the governing field equations. This step is sometimes called hole-cutting or grid-point blanking, and is the first step in the domain connectivity (or grid assembly) process. At grid boundaries where a flow solver boundary condition is not specified, and at the boundaries of holes from the hole-cutting process, the solution needs to be interpolated from neighboring overlapping grids. The second step in domain connectivity involves the search for donor cells (interpolation stencils) for the fringe points on such boundaries.

Many algorithms and software have been developed over the years to perform hole-cutting. Some of the more popular ones in use today include a Cartesian hole map and line-of-sight method in PEGASUS^{5,7}, implicit hole-cutting in SUGGAR++,⁸ PUNDIT,⁹ and OVERTURE,¹⁰ and X-ray hole-cutting¹¹ in the DCF module of the OVERFLOW flow solver (referred to as OVERFLOW-DCF from here on). These schemes and software have all been successfully employed on many applications. In problems involving components in relative motion, hole-cutting needs to be performed at every time step. The X-rays approach is one of the fastest available in this regime since hole-cutting involves only a direct Cartesian map look-up. The Cartesian map is generated only once at the beginning for rigid-body motion and is transformed with the grids as they move relative to each other. Additionally, the number of grid points requiring donor stencil search in the X-rays approach is typically much smaller compared to that required for implicit hole-cutting schemes. For the X-rays approach, only L_f surface layers of fringe points require stencil search where L_f is

*Computer Scientist, AIAA Senior Member

†Aerospace Engineer, AIAA Senior Member

‡Aerospace Engineer, AIAA Associate Fellow

This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States.2013

the number of fringe layers imposed at the outer boundaries and hole boundaries of a grid. For implicit hole-cutting schemes, every grid point in the entire volume grid system requires a donor stencil search. Hence, the number of stencil searches for the X-rays approach is proportional to the number of points on a surface layer, while the number of stencil searches for the implicit hole-cutting approach is proportional to the number of points in a volume domain. Therefore, for the same stencil-search algorithm, the total stencil search time for the X-rays scheme is much faster than that for the implicit hole-cutting method simply because less items are sent to the search routine.

Although the X-rays approach has a clear advantage in computational speed for moving-body problems, it has drawbacks in the setup process where the inputs typically require significant user effort and expertise. This problem has been addressed in Ref. 12 which presented methods to automate all of the inputs required. The only inputs needed are the flow solver boundary conditions for each grid, and a component ID that is assigned to each solid wall surface in the grid system. The scheme presented in Ref. 12 is implemented into a software tool called C3P which uses an adaptive X-rays method to create a minimum hole. This leaves the hole boundary of coarser neighboring grids immediately next to fine viscous cells adjacent to solid walls. The next step involves using an auxiliary wall-distance function (see Section 2 below) to provide a first estimate of the hole boundary away from the walls, followed by iterations of the hole boundary to remove orphan points (grid points that do not have a donor cell due to insufficient overlap). While this scheme performed well, the auxiliary wall distance function was expensive to compute. The first objective of the current work is to investigate faster ways to determine the auxiliary wall-distance function.

In the original X-rays method, the hole boundary offset from the minimum hole is specified as a constant distance by the user. In the new approach,¹² a spatially-varying hole boundary offset is automatically generated. Various criteria have been used in different software to determine the location of the hole boundary. In PEGASUS5 and SUGGAR++, a cell compatibility measure between the cell containing the fringe point and its donor cell is utilized. Examples of this measure include cell volume, aspect ratio, orientation, maximum length of cell diagonal, and any combination of the above. Applications of different measures result in different locations of the hole boundaries, all of which may be considered to be valid to the extent that there are no orphan points present. However, it is unclear which measure provides the most desirable results such as accurate aerodynamic loads and fast solution convergence. The second objective of this work involves an investigation in the variations in converged values and convergence rates of aerodynamic loads for different hole boundary locations and different values of cell attribute compatibility. The study is performed using four test cases: an isolated re-entry capsule, a two-rocket configuration, the Common Research Model from the AIAA 4th Drag Prediction Workshop,¹³ and the D8 “Double-Bubble” Aircraft.¹⁴ Each of these cases was run using three different domain connectivity codes: standard X-rays in OVERFLOW-DCF, new enhanced X-rays in the C3P software, and the PEGASUS5 overset software. The flow solutions for all cases were computed using the OVERFLOW flow solver.¹⁷

II. Hole Boundary Estimate Using an Auxiliary Wall-Distance Function

The hole boundaries obtained after the minimum hole-cut step are usually not accurate for information transfer between grids since fine viscous grid cells near the wall are serving as donor cells for fringe points from much coarser regions of neighboring grids. In Ref. 12, a first estimate of a more appropriate hole boundary location away from the minimum hole is provided by a wall-distance function. It was found that in order to expand from the minimum hole on the surfaces of intersecting components such as a wing-body junction, the standard wall-distance function used for solving turbulence model equations is insufficient. A more sophisticated auxiliary wall-distance function that is sensitive to distances between components down to the wall surface is needed. Figure 1 shows such a function where the distance between intersecting components is small (blue) at the wing-body junction, and similarly for the tail-body junction.

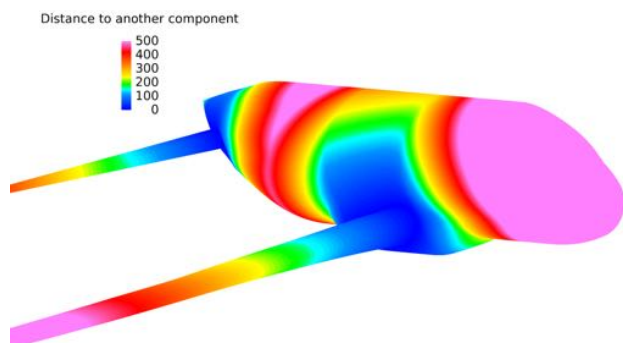


Figure 1. Auxiliary wall-distance function on surface of wing-body-tail configuration with small values colored blue and large values colored magenta.

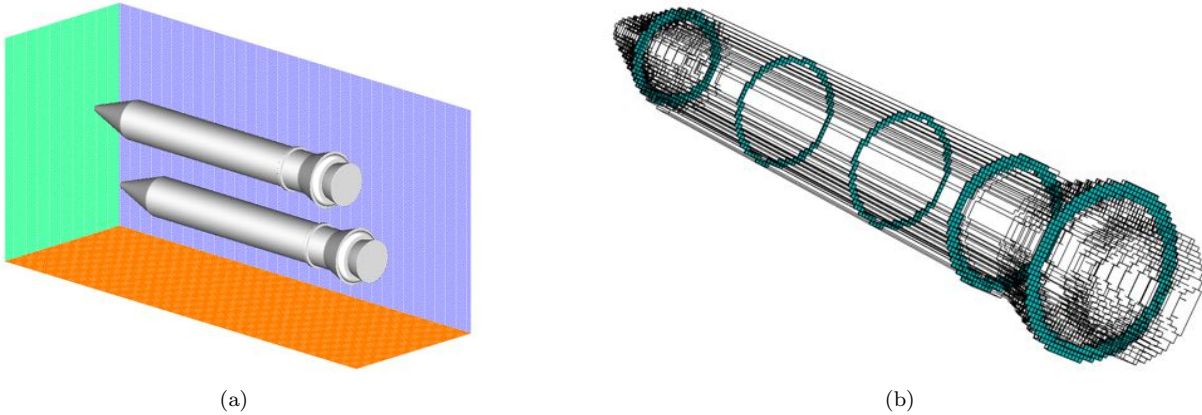


Figure 2. Reference Cartesian box grid for auxiliary wall-distance function computation on two-rocket test case. (a) Volume grid slices for reference Cartesian box and rockets surface geometry. (b) Cells on reference Cartesian box intersected by surface cells on one of the rockets.

The automated hole-cutting scheme in Ref. 12 begins by identifying the geometric components in the grid system. Let the total number of components in the system be N_c . Each near-body curvilinear grid point in the volume grid system is associated with a component by tracing a grid line from the grid point to the wall surface of the component. Grid points in off-body grids that do not contain a wall are not associated with any component. For convenience, these grid points are defined to be associated with component ‘zero’.

The first hole boundary estimate from the minimum hole requires the computation of the distance from each grid point to the closest wall belonging to each component that is not associated with the point. Suppose a grid point is associated with component M . This scheme requires the computation of the distance from the point to each component other than M . In order to determine the hole boundary estimate for all grid points in the system, N_c wall-distance function computations are needed over the entire grid domain. For this new approach to be competitive in speed compared to the original X-rays scheme, a fast method to compute this auxiliary wall-distance function is needed. One such scheme was presented in Ref. 12 but was found to be too slow. A very different but much faster scheme is now outlined below. It is noted that for purpose of hole-cutting, an approximate wall-distance is sufficient to provide satisfactory results.

The first step involves constructing a uniform Cartesian reference box grid around all near-body volume grids in the configuration (Fig. 2a). These are volume grids that contain a solid wall boundary. The uniform spacing for this reference Cartesian grid is chosen to be the average surface grid spacing of all near-body grids. A user-specified maximum (e.g., 10 million points) is used to limit the total number of grid points in this Cartesian grid.

For a given component i , the Cartesian grid cells that intersect the component surface are quickly identified (Fig. 2). These intersecting-cells are easily determined by finding the bounding box of each triangle on a triangular tessellation of the component surface, and then marking the reference Cartesian grid cells that are intersected by the bounding box (Fig. 2b). For each vertex on the intersecting-cells, an accurate computation of the distance to the component surface is performed. The wall-distance function values on all neighboring vertices of the intersecting-cell vertices are approximated by adding the known values of the wall-distance function to the Cartesian cell size. This marching scheme is repeated until an approximate wall-distance has been determined for all remaining Cartesian grid vertices that do not belong to the original intersecting-cells.¹⁵ After the fast marching scheme has been applied to all components, a wall-distance function W_i is available on the reference Cartesian grid for each component i .

Given any volume grid point P that is associated with component M , the distance to a wall that does not belong to M is quickly determined by a Cartesian cell look-up in the reference Cartesian grid, followed by interpolation of the approximate wall distance to P . For a volume grid point P that is not associated with any component (one from an off-body grid with no solid wall boundary), the same direct look-up procedure can be used to determine its approximate wall distance if the point lies inside the reference Cartesian grid.

If the point P lies outside the reference Cartesian grid as shown in Figure 3, its distance D_Q to the closest point Q on the outer boundary of the reference Cartesian box can be easily computed. The approximate wall distance for this point is then the sum of D_Q and the approximate wall distance at Q .

Figure 4a shows the auxiliary wall-distance function on a slice of the reference Cartesian box grid for component ‘zero’ which is associated with all grid points belonging to any off-body grid (a grid that does not contain a wall). It can be seen that this function is identical to the standard wall-distance function, i.e., finding the minimum distance to all walls in the system (distance to both rockets). Figure 4b shows a similar function for the top rocket only. This is used to mark up points in the bottom rocket that fall too close to the top rocket. Such points are to be converted from field to hole points.

After an initial estimate of an offset hole boundary using the auxiliary wall-distance function, some orphan points may be present due to insufficient grid overlap. Ref. 12 describes a hole boundary adjustment procedure to reduce the number of orphan points. If implemented correctly, this procedure should produce no more orphan points compared to the minimum hole-cut state. Figures 5a,b show the final adjusted hole boundaries for the Trapwing 3-element high-lift configuration.¹⁶ The automated procedure is able to create a spatially-varying hole-boundary offset distance in the tight gap between the high-lift elements. Moreover, inter-grid communication occurs mostly between cells of comparable sizes. In the standard X-rays approach, the hole boundary would have been limited to a small constant offset distance which would result in large discrepancies in the cell sizes of fringe points and their donor stencils in some regions. Creating a variable hole-boundary offset distance using standard X-rays is possible, but extra manual effort would be required to construct additional hole-cutting surfaces.

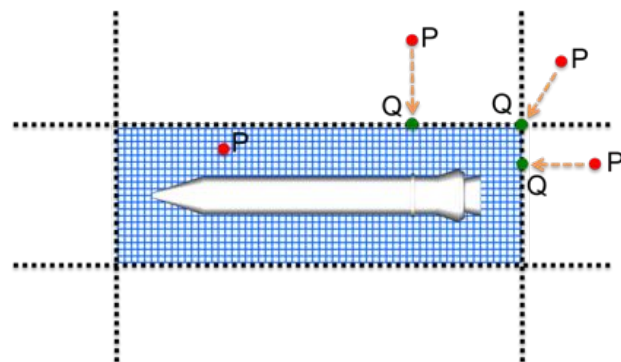


Figure 3. Computation of approximate closest distance from point P to the wall. If P is inside the reference Cartesian grid, direct interpolation is used. If P is outside the reference Cartesian grid, the distance from P to point Q is added to the wall distance at Q , where Q is the closest point on the outer boundary of the reference Cartesian mesh from P .

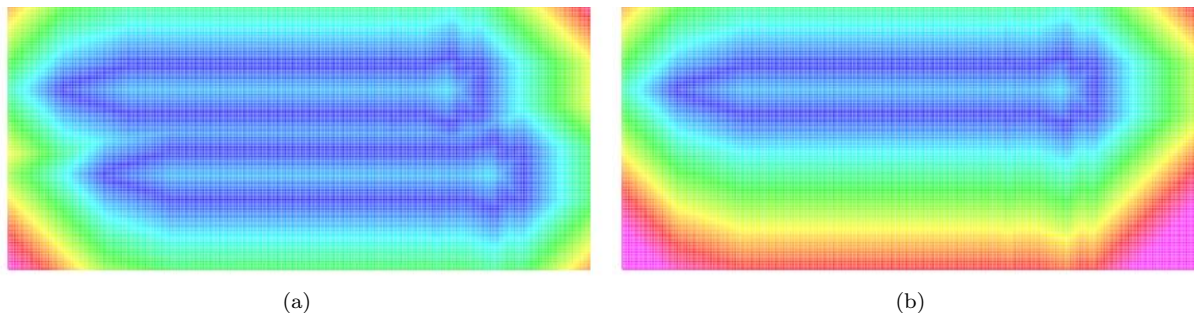


Figure 4. Approximate auxiliary wall-distance function on volume grid slice of reference Cartesian box for two-rockets test case. (a) Distance to both rockets. (b) Distance to top rocket only.

Table 1. Comparison of wall-clock time (minutes) for various test cases on Xeon workstation with Intel Fortran compiler for running domain connectivity using OVERFLOW-DCF (original X-rays) with 8 MPI ranks, C3P (improved X-rays) with 8 OpenMP threads, and PEGASUS5 with 8 MPI ranks (XV-15 = rotorcraft with hub and blades, DPW4 = AIAA 4th Drag Prediction Workshop model, 2-Rockets = two-rocket configuration, HLWS1 = AIAA 1st High Lift Prediction Workshop 3-element wing configuration (Trap Wing),¹⁶ D8 = D8 Double Bubble Aircraft).

Cases	# Grid points ($\times 10^6$)	DCF	C3P	PEG5	% change in C3P from DCF
XV-15	10.0	0.3	0.6	0.75	+100%
DPW4	16.8	1.2	1.5	1.7	+25%
2-Rockets	24.5	0.5	0.8	1.0	+60%
HLPW1	50.6	2.3	1.7	4.2	-26%
D8	83.0	11.5	7.2	8.2	-37%

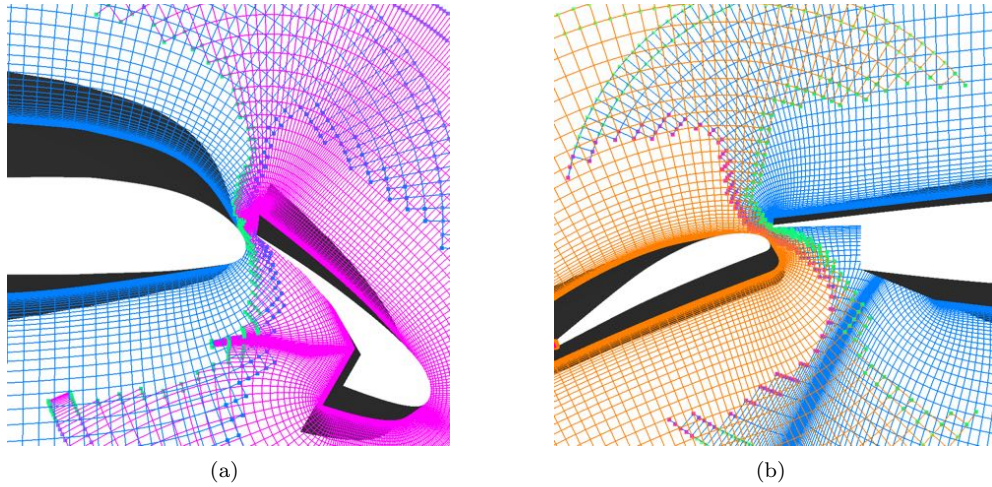


Figure 5. Hole boundaries for the Trapwing 3-element high-lift configuration with fringe points marked by symbols. (a) Gap between slat and wing. (b) Gap between wing and flap.

The new auxiliary wall-distance function computation scheme has been implemented into the domain connectivity software C3P.¹² A factor of 2.5 to 12 times speed up was observed in different cases in using the new wall-distance computation algorithm described above compared to the scheme presented in Ref. 12. It is anticipated that the total domain connectivity computational time (hole-cutting and donor stencil search) for the new approach in C3P should be more than that for the original X-rays approach in OVERFLOW-DCF. In the original scheme, the X-ray maps and offset distances from the minimum hole are manual prescribed, and donor stencil search is performed just once for all the identified fringe points. In the present method implemented in C3P, extra work is required to automatically build the adaptive X-ray maps, compute the auxiliary distance function, iterate on the hole boundaries to remove the orphan points, and perform donor stencil search multiple times during the hole boundary iterations. Despite this extra work, it is found the computational cost for C3P is not that much more than the original scheme in OVERFLOW-DCF, and usually less than the scheme utilized in PEGASUS5. This is illustrated by the wall clock times presented in Table 1 for running five test cases on a Xeon desktop machine using the Intel Fortran compiler. The results show that the new scheme is slower than the efficient original X-rays scheme for the smaller test cases, but by no more than a factor of two. For larger test cases, the new scheme appears to perform better than the original. Further testing on a wider variety of cases is needed to determine if the observed timing trends above hold in general.

It is also noted that for OVERFLOW-DCF and C3P, the domain connectivity parallel efficiency is typically around 25% - 40%, while that for PEGASUS5 is typically around 65% - 80%. This is because OVERFLOW-DCF's load balancing scheme has been constructed to optimize on flow solver efficiency and not on domain connectivity efficiency. No attempt has been made to perform load balancing on the computational work required for C3P. Even with a factor of two or three slower in CPU wall-clock time in using C3P compared to the original X-rays scheme, savings in manual effort time with C3P typically far outweighs any CPU time disadvantages.

III. Aerodynamic Loads Sensitivity Study

For many years, users of overset grid methods have used various algorithms and software to create hole-cuts on complex grid systems. As long as the hole boundaries have been displaced away from the minimum hole, and the number of orphan points has been reduced to zero or a reasonable number ($<0.01\%$ of the total number of grid points in the system), the grid system is deemed acceptable for flow solution computation. Unlike the minimum hole which can be uniquely defined by tagging grid points that fall inside the solid boundaries of components, the offset or final hole boundary used for flow computations is not uniquely defined. In algorithms such as the original X-rays approach, the final hole boundary is defined by the user via a constant distance offset from the minimum hole. In more complex algorithms, the final hole boundary is automatically generated based on matching cell attributes between fringe points and their interpolation donor cells. The hole boundaries tend to be smooth in the former, while they tend to be more ragged in

the latter. In both cases, the hole boundaries are away from the minimum hole and orphan points may be absent.

It has long been adopted as best practice that grid communication should occur between grid cells of comparable attributes so that flow gradients resolved by one grid can be accurately transferred to its neighboring grid. However, there has not been a consensus on which grid cell attribute, or combination of attributes, should be used for controlling the hole boundary location. Attributes in use today include ratios of cell volume, aspect ratio, orientation, maximum cell diagonal length, and others.^{7,8} Furthermore, to the authors' knowledge, the effects of utilizing different attributes on the converged value and convergence rates of aerodynamic loads have not been studied. The work presented in this section is an investigation into such effects.

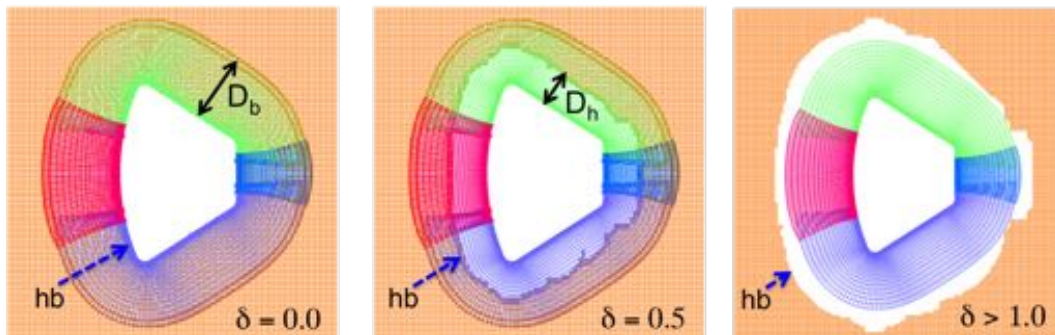


Figure 6. Normalized hole boundary offset between near-body and off-body grids (D_b = approximate normal distance to outer boundary for near-body grid, D_h = hole boundary offset distance, δ = normalized hole boundary offset distance, hb indicates hole boundary location of off-body Cartesian grid). (a) $\delta = 0$. (b) $\delta = 0.5$. (c) $\delta > 1$.

In order to provide a fair comparison between different test cases, a normalized hole boundary location or offset is defined. Three situations can arise in hole boundaries between grids.

1. A component cutting a hole in an off-body grid (Fig. 6).
2. Disjoint components cutting holes in grid points associated with other components (Fig. 7a).
3. A component cutting a hole in grid points on and close to the surface of another intersecting component (Fig. 7b).

In case (1), let D_b be the approximate normal distance from the surface to the outer boundary of the near-body grid(s) associated with the component. Also let D_h be the hole boundary offset distance from the minimum hole. Then, a normalized hole boundary offset distance δ can be defined to be

$$\delta = D_h / D_b \quad (1)$$

Figure 6 illustrates the hole boundary locations for various values of δ for case (1).

In case (2), let D_m be the minimum distance between two disjoint components (Figure 7a). Then, a normalized hole boundary offset distance δ between the components can be defined to be

$$\delta = D_h / (0.5D_m) \quad (2)$$

In this case, the denominator on the right-hand-side is half of the minimum distance between the components. If a hole boundary offset D_h is created to lie at $> 0.5D_m$, there will be no overlap between the grids from the two components at the minimum distance location. The situation at $D_h = 0.5D_m$ should then correspond to a normalized offset distance of $\delta = 1$.

In case (3), let D_c be the surface grid marching distance of the collar grid between two intersecting components (Figure 7b). Then, a normalized hole boundary offset distance δ between the components can be defined to be

$$\delta = D_h / D_c \quad (3)$$

Similar to case (2), a hole boundary offset $D_h > D_c$ would result in no overlap between the collar grid and the surface grids associated with the intersecting components. The situation at $D_h = D_c$ would correspond to a normalized offset distance of $\delta = 1$.

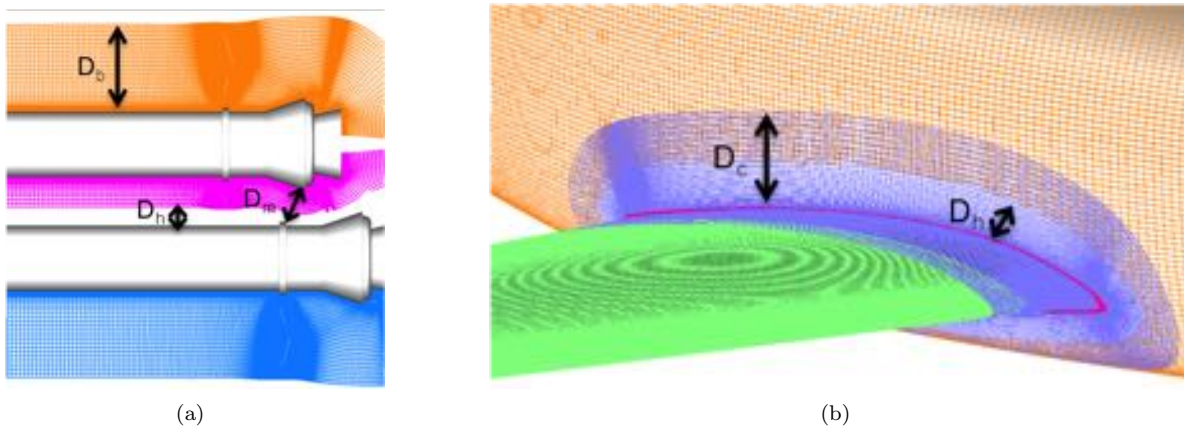


Figure 7. Normalized hole boundary offset distance δ for disjoint and intersecting components. (a) Two disjoint components with $\delta = D_h / (0.5 D_m)$. (b) Two intersecting components with $\delta = D_h / D_c$. D_h = hole boundary offset distance, D_m = minimum separation distance between disjoint components, D_c = collar grid surface marching distance between intersecting components.

In the current study, the cell volume is used as the compatibility measure between a fringe point and its donor cell. No attempt is made to determine which attribute provides the best results. The cell volume is selected as a simple measure to illustrate the differences that can be obtained in the behavior of the converged values and convergence rates of the aerodynamic loads.

Let V_f and V_i be the cell volumes at a fringe point and its donor cell, respectively. The compatibility c_i is defined to be

$$c_i = \min(V_f/V_i, V_i/V_f) \quad (4)$$

A value of c_i close to one means the cell volume at the fringe point and volume of its donor cell are very similar, while low values of c_i means the corresponding cell volumes are vastly different. Orphan points are given a value of zero for c_i . By averaging c_i over all fringe points, a global measure of average compatibility C can be computed.

For the test cases presented in the subsequent subsections, three methods for determining hole boundaries are compared using three software packages: standard X-rays in OVERFLOW-DCF, enhanced X-rays in C3P, and the PEGASUS5 software. In standard X-rays, a constant user-specified offset distance is used to create the final hole boundary from the minimum hole. In C3P, a wall-distance function is used to obtain an initial offset of the hole boundary from the minimum hole. Iterations on the hole boundary then follow to drive the number of orphan points to zero. In PEGASUS5, grid points not in the minimum hole go through an implicit hole-cut process. These grid points are either left as field points or marked as fringe points based on a cell attribute parameter which includes cell volume, aspect ratio, and orientation.

The OVERFLOW¹⁷ compressible viscous flow solver is used to compute the flow solutions for all test cases. Since the current study is concerned with investigating the differences in convergence behavior between cases with different hole boundary locations, it is not critical that the cases be run to a converged steady state. The absolute values of the converged loads is not important, but the runs were made to continue well past the initial transient until a clear trend was observed between cases with different hole boundary locations. The aerodynamic load convergence history and final values are presented for normalized hole boundary offset distances $\delta = 0.0, 0.1, 0.25, 0.5, 0.75, 0.9$ using standard X-rays, and for the automatically generated hole boundaries from the C3P and PEGASUS5 software. For each case, the average cell size compatibility C between fringe and donor cells is also reported for the hole-boundary fringe points.

III.A. Isolated Re-entry Capsule

The isolated capsule test case consists of three near-body grids around the capsule embedded in a Cartesian box grid with a total of 11 million grid points in the system. Uniform spacing is used in the Cartesian core that matches the near-body grid outer boundary spacings, and stretched spacing is used to extend the Cartesian grid to the far field. The solution is computed at a free stream Mach number of 1.2, angle of attack of 180 degrees (flow directed at the heat shield), and a Reynolds number per inch of 18200.

Figure 8a shows the automatically created hole boundary by C3P for this single-component test case. In this simple case, the hole boundary location is almost identical to the manually specified $\delta = 0.5$ case.

Figure 8b shows the Mach number contours for the converged solution with a bow shock in front of the heat shield. No orphan points are present for the C3P and PEGASUS5 cases, and for all δ cases except for $\delta = 0.9$ where there are 930 orphan points.

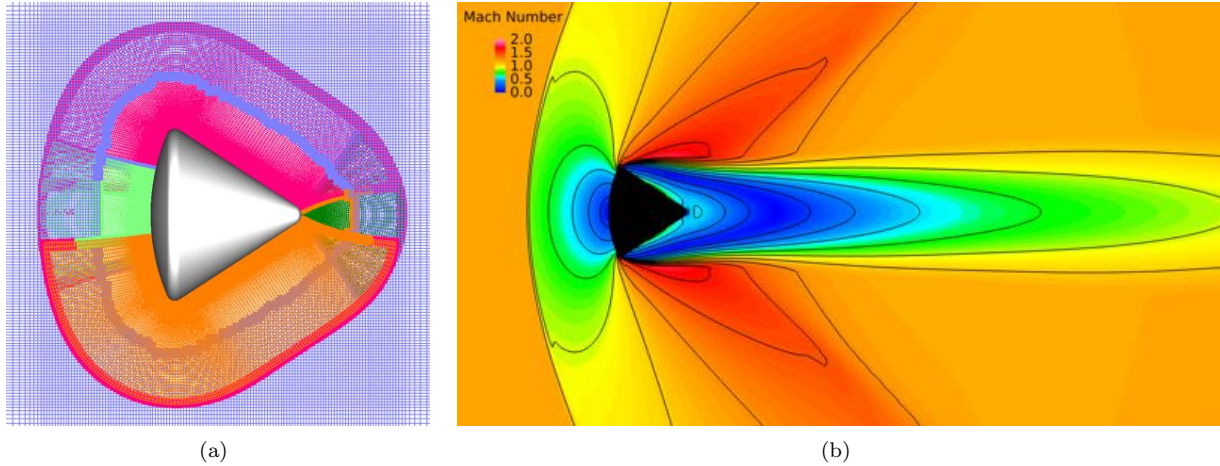


Figure 8. Isolated re-entry capsule test case. (a) Volume grid slices for near-body grids and off-body Cartesian grid showing hole boundary location created automatically by C3P. Fringe points are marked by symbols. (b) Mach number contours.

With an angle of attack of 180 degrees, the lift and pitching moment coefficients remain practically zero as expected. Figure 9a shows the convergence histories of the drag coefficient for the different cases. All cases appear to converge at about the same rate but to slightly different converged values. The C3P (cyan) result falls close to the $\delta = 0.5$ case (blue) as expected since the hole boundaries of the two cases are almost identical. On either side of the $\delta = 0.5$ case are the $\delta = 0.25$ and 0.75 cases. The PEGASUS5 and $\delta = 0.9$ cases have converged C_D 's furthest from the mean of the test cases (about 0.6% deviation from the mean).

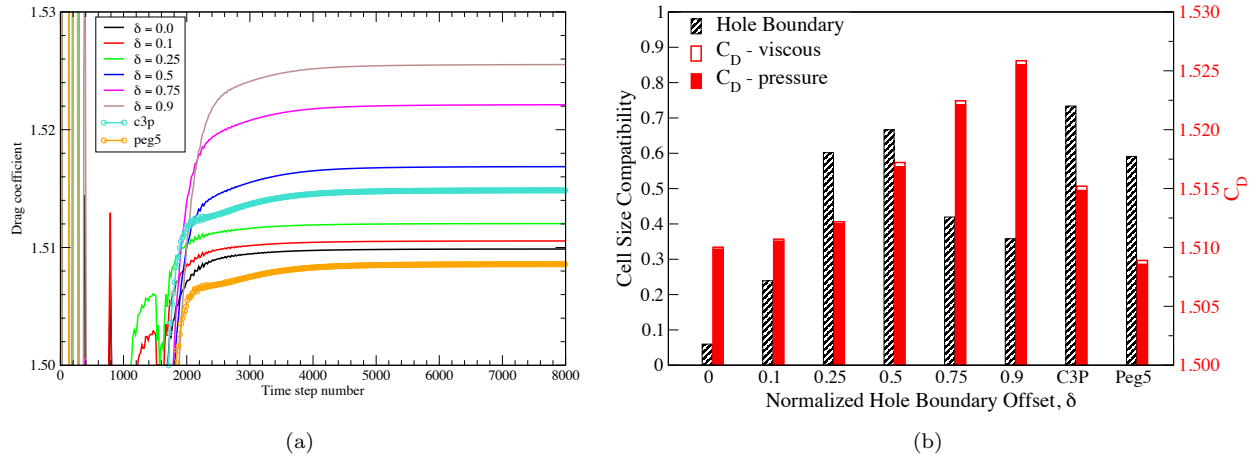


Figure 9. Isolated re-entry capsule test case drag coefficient convergence and compatibility. (a) Drag coefficient C_D convergence history for various δ , C3P, and PEGASUS5. (b) Hole boundary fringe points interpolation compatibility and converged value of C_D for various δ , C3P, and PEGASUS5.

Figure 9b shows the average cell size compatibility between fringe points and donor cells for hole boundary points (black and white stripe bars) for various hole boundary locations and treatments. As can be seen from the relative grid resolutions of the volume grid slices in Figure 8a, the $\delta = 0.5$ case has the best compatibility among the OVERFLOW-DCF cases for the hole boundary, with the compatibility falling lower as δ deviates from 0.5. However, the final value of C_D does not appear to correlate well with averaged cell size compatibility. For example, the cell size compatibility for the $\delta = 0.25$ and PEGASUS5 cases are almost the same, but the final C_D values are quite different.

Figure 10 shows the variation of the integrated sectional drag against streamwise distance x . The plot for all cases fall practically on top of each other. Only small differences are observed in the flow separation

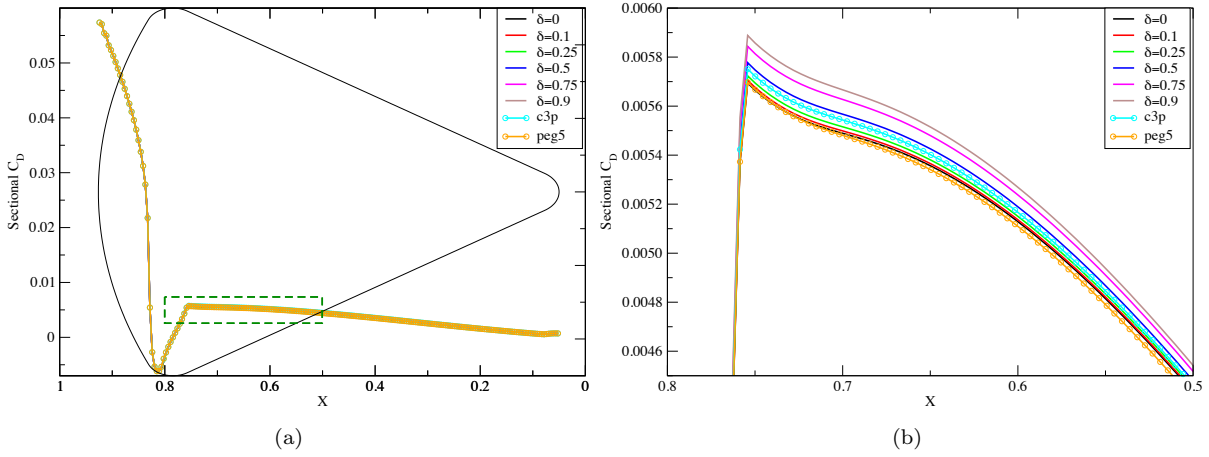


Figure 10. Isolated re-entry capsule test case sectional drag coefficient C_D for various δ , C3P, and PEGASUS5. (a) Overall view. (b) Zoomed view.

region behind the shoulder of the capsule. It is well known that this case, with an angle of attack of 180 degrees, is highly sensitive in the shoulder region. Small perturbations in the numerical scheme, such as hole boundary location, could trigger slight differences in the flow, resulting in small differences in the flow separation pattern.

III.B. Two-Rockets Configuration

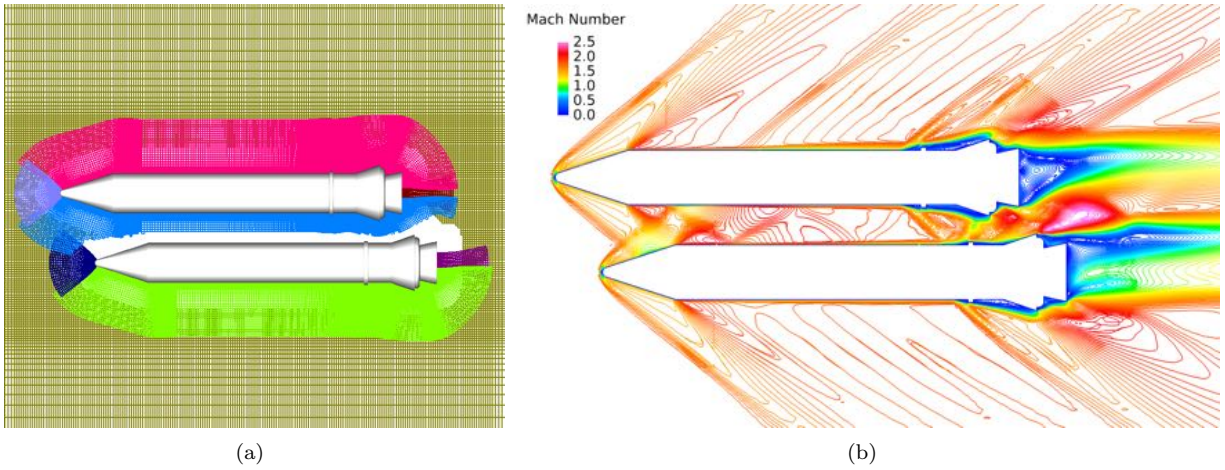


Figure 11. Two-rockets test case. (a) Volume grid slices for near-body grids and off-body Cartesian grid showing hole boundary location created automatically by C3P. (b) Mach number contours.

The two-rockets test case consists of two identical rockets with three near-body grids each embedded in a stretched Cartesian box grid with a uniform core region. The entire grid system contains 7 grids and 24.5 million grid points. An offset exists between the two rockets in both the axial and transverse directions. The solution is run at a freestream Mach number of 1.8, with angle of attack at 0 degree, and a Reynolds number per inch of 2×10^5 .

Table 2. Number of orphan points for various δ 's, C3P, and PEGASUS5 for two-rockets test case.

Hole boundary treatment	$\delta = 0.0$	$\delta = 0.1$	$\delta = 0.25$	$\delta = 0.5$	$\delta = 0.75$	$\delta = 0.9$	C3P	PEG5
No. of orphan points	0	0	12	195	1722	75622	0	0

Figure 11a shows the automatically created hole boundary from C3P for the upper rocket in this two-disjoint-components test case. A spatially-variable hole-boundary offset is created in the near-body grids

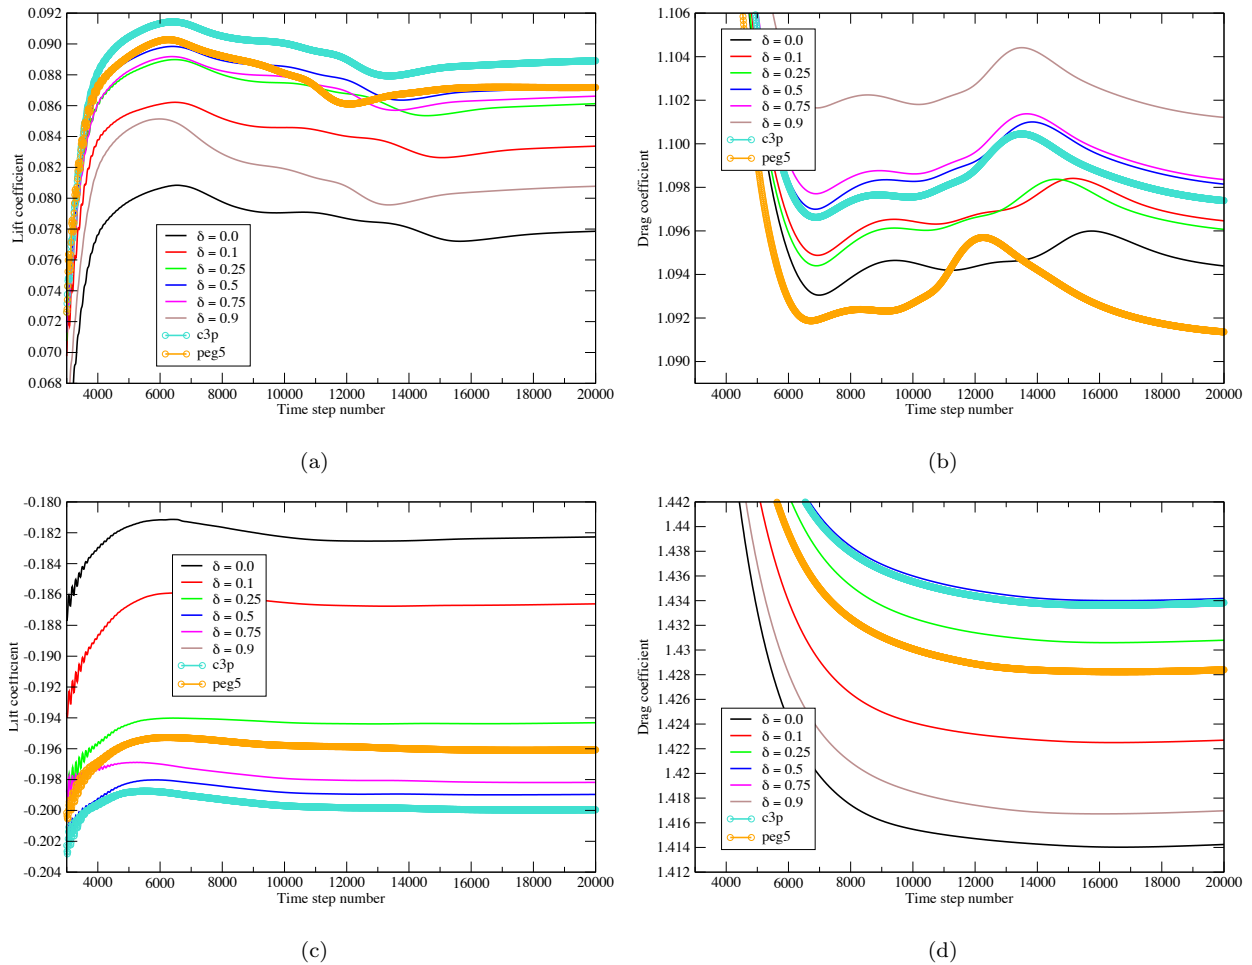


Figure 12. Two-rockets test case lift (C_L) and drag (C_D) coefficient convergence histories. (a) Top rocket C_L . (b) Top rocket C_D . (c) Bottom rocket C_L . (d) Bottom rocket C_D .

between the two rockets. Figure 11b shows the Mach number contours for the converged solution with complex shock reflections in the region between the rockets. The number of orphan points for different δ 's, C3P, and PEGASUS5 are shown in Table 2. All cases have either zero or a small number of orphan points ($<0.01\%$) relative to the total number of grid points, except for the $\delta = 0.9$ case whose number of orphan points is about 0.3% of the total number of grid points.

Figure 12 shows the convergence histories of the lift (C_L) and drag (C_D) coefficients for the top and bottom rockets. The cases appear to converge at about the same rate for different δ 's, but to different final values as in the isolated capsule case. For C_L of both rockets and for C_D of the bottom rocket, results from C3P, PEGASUS5 and the 'best practice' values of δ 's (0.25, 0.5, 0.75) form a closely packed group. For C_D of the top rocket, no clear trend is observed.

Figure 13 shows the average cell size compatibility measure between hole boundary fringe points and donor cells for various δ 's, C3P, and PEGASUS5. The $\delta = 0.75$ case appears to have the best compatibility value but the final converged loads do not appear to vary much between all the cases except for the $\delta = 0$, 0.1, and 0.9 cases.

Figure 14 shows the surface pressure coefficient profile on a cut through the $y = 0$ symmetry plane for the bottom rocket. Such a cut produces two curves, one along the side facing towards the top rocket ($+z$ side), and one along the side facing away from the top rocket ($-z$ side). The plots demonstrate that there is very little difference between the $\delta = 0.1$ case, and the C3P and PEGASUS5 cases. The zoomed views in Figure 14 show that no consistent conclusions can be drawn about the small differences between the three cases.

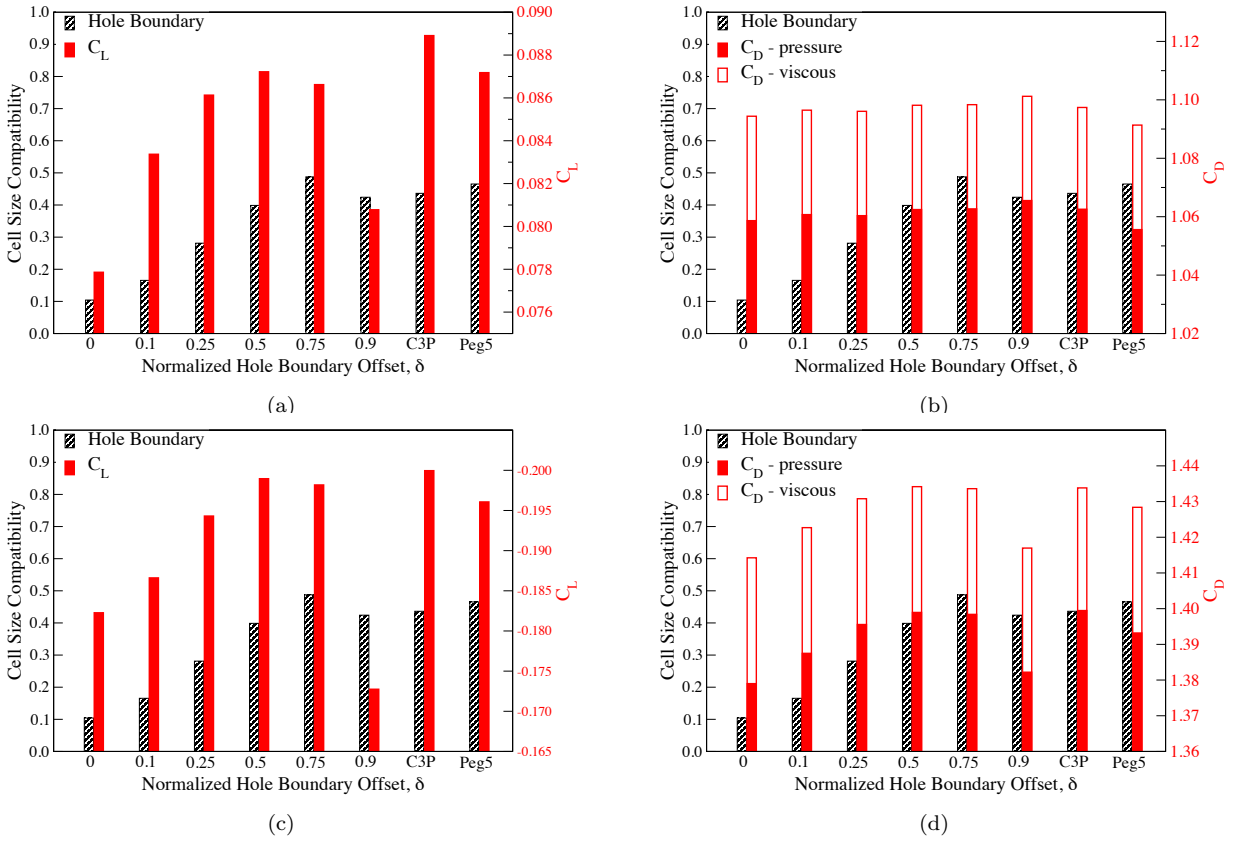


Figure 13. Two-rockets test case hole boundary fringe points interpolation compatibility. (a) Top rocket C_L . (b) Top rocket C_D . (c) Bottom rocket C_L . (d) Bottom rocket C_D .

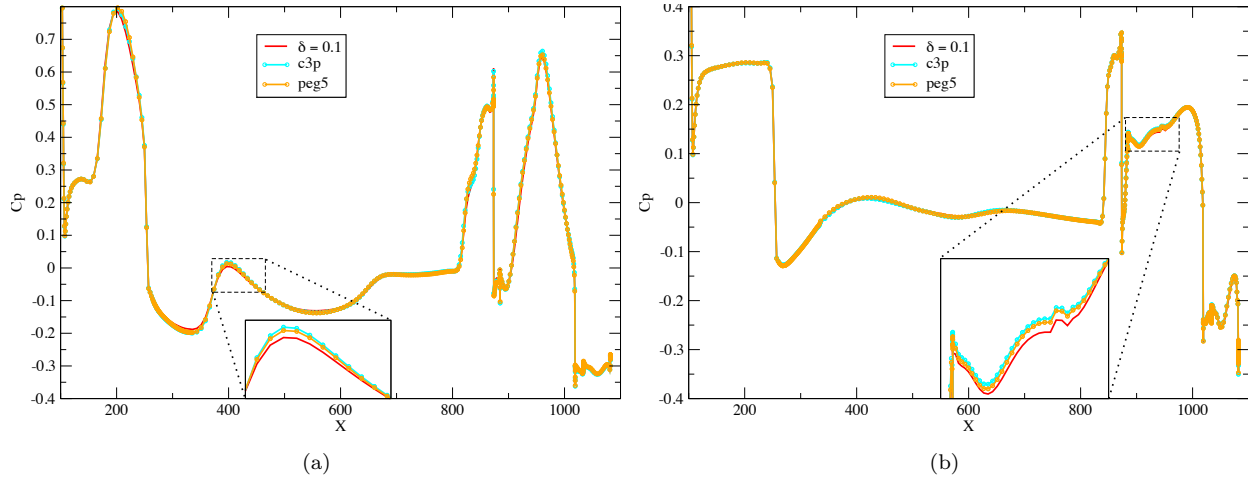


Figure 14. Pressure coefficient against streamwise distance at $y = 0$ symmetry plane cut through bottom rocket. (a) Side facing towards top rocket. (b) Side facing away from top rocket.

III.C. AIAA 4th Drag Prediction Workshop Common Research Model

The DPW4 test case is the Common Research Model (CRM) from the 4th AIAA Drag Prediction Workshop. It is a half-body configuration with three components: a fuselage, a wing and a horizontal tail. The grid system includes near-body grids for each of the three components, a collar grid between the fuselage and wing, a collar grid between the fuselage and tail, and a stretched Cartesian box grid with a uniform core region. The entire grid system contains 17 grids and 16.8 million grid points. The solution is run at a freestream Mach number of 0.85, an angle of attack of 2.363 degrees, and a Reynolds number per inch of 18129.

The grid system downloaded from the DPW4 website contains a C-grid around the wing with a one point collapse over the finite-thickness trailing edge, and a cap grid that wraps around the trailing edge with 25 points over the trailing edge face (Figure 15). In the test runs, three different treatments of this trailing edge region were compared. For the standard X-ray runs, and for the first C3P run, the C-grid points were left as field points (Figure 15a). For the second C3P run, the C-grid points were explicitly blanked in the vicinity of the trailing edge (Figure 15b). In the PEGASUS5 case, no explicit hole-cutting of the C-grid points was done, but PEGASUS5 added extra layers of fringes to create a virtual-hole and overlap using its cell-difference algorithm (Figures 15c, d). In the second C3P and in the PEGASUS5 cases, the cap grid is the only grid responsible for resolving the flow in the trailing edge region.

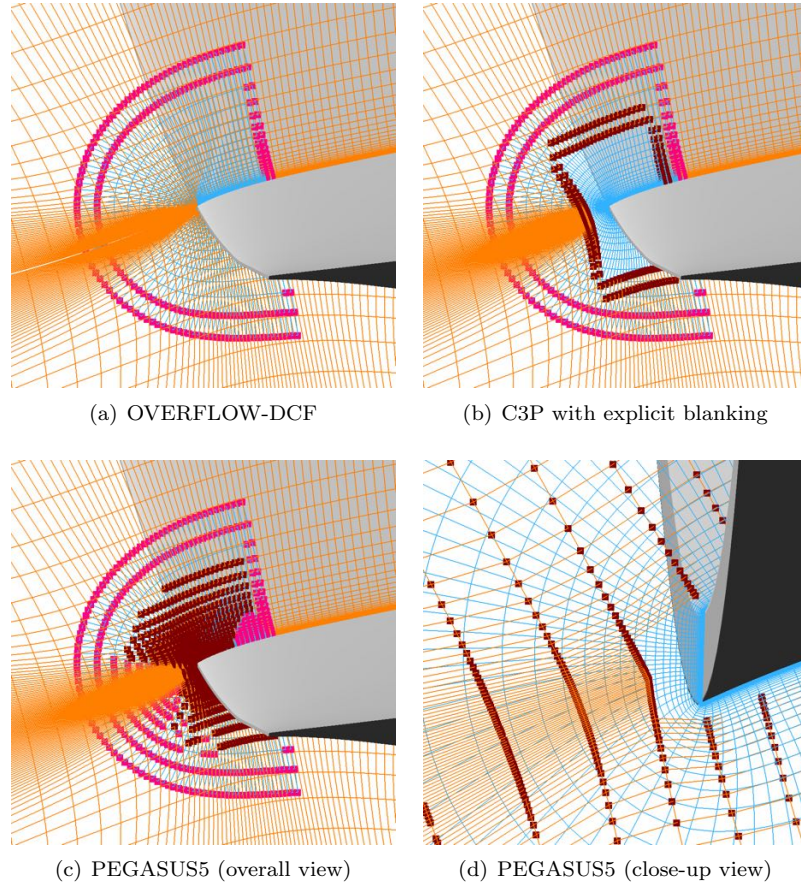


Figure 15. Trailing edge treatment of fringe points and iblanks for DPW4 test case. Fringe points are marked by solid symbols. (a) OVERFLOW-DCF: C-grid points treated as field points. (b) C3P: C-grid points treated as hole points. (c) PEGASUS5: C-grid points treated as second-level fringe points. (d) PEGASUS5: close-up view near trailing edge surface.

Figure 16a shows the automatically created hole boundary by C3P for this intersecting components test case. Figure 16b shows the pressure coefficient contours for the converged solution with a shock over the wing. The number of orphan points for different δ 's, C3P, and PEGASUS5 is shown in Table 3. All cases have either zero or a small number of orphan points ($<0.04\%$) relative to the total number of grid points except for the $\delta = 0.9$ case (0.2%).

Table 3. Number of orphan points for various δ 's, C3P, and PEGASUS5 for DPW4 test case.

Hole boundary treatment	$\delta = 0.0$	$\delta = 0.1$	$\delta = 0.25$	$\delta = 0.5$	$\delta = 0.75$	$\delta = 0.9$	C3P	PEG5
No. of orphan points	0	0	0	3027	6017	37546	0	0

Figures 17a, 17c, 17e show the convergence histories of the lift (C_L), drag (C_D), and pitching moment (C_M) coefficients. The cases appear to converge at about the same rate for different δ 's and C3P, with the case from PEGASUS5 being more oscillatory. Final values of the loads appear to fall into two groups. In

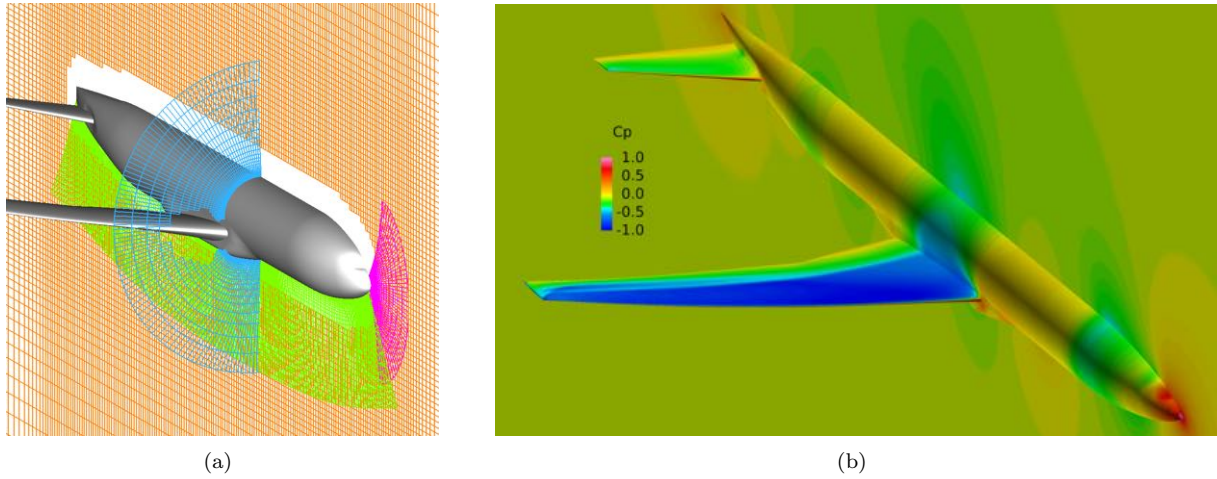


Figure 16. DPW4 test case. (a) Volume grid slices for near-body grids and off-body Cartesian grid showing hole boundary location created automatically by C3P. (b) Pressure coefficient contours.

group 1, the influence of the C-grid is removed in the trailing edge region by blanking (C3P) or by marking grid points as second-level fringe (PEGASUS5). In group 2, the C-grid points are left as field points in the trailing edge region. Figure 17a shows that the group 1 cases produce a final lift coefficient that matches closer to the experimental value of 0.5.

The maximum deviation of the final value of any member of the first group to the mean of the group is about 3% for C_L , 1.7% for C_D (less than 0.5 drag count), 1.6% for C_M . The maximum deviation of final value of any member of the second group (excluding $\delta = 0.0$ and 0.9) to the mean of the group is about 0.07% for C_L , 2.8% for C_D (less than 1 drag count), and 3.3% for C_M . From the DPW4 workshop, many structured grid flow solvers using different numerical schemes and turbulence models also reported converged values of C_D and C_M . The maximum deviation from the mean for this group is about 7 drag counts for C_D and 18.9% for C_M , which is much larger than the differences observed for different hole boundary locations in the present study.

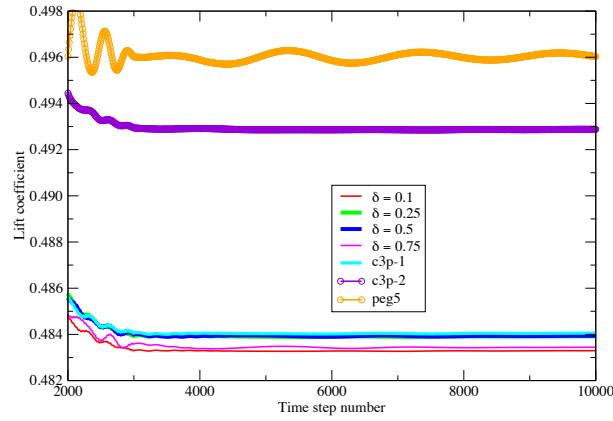
Figures 17b, 17d, 17f show the cell size compatibility between the hole boundary fringe points and donor cells for various δ 's, C3P, and PEGASUS5. Cell size compatibility for hole boundary points are best for C3P and PEGASUS. However, the final load values are similar for cases with different compatibility values except for the $\delta = 0.0$ and 0.9 cases.

In Figure 18a, the plot for sectional force coefficient in z against spanwise distance y shows that variations between the different cases occur primarily over the wing, with almost no variation over the fuselage. On the plot of the same force coefficient against streamwise distance x as depicted in Figure 18b, the same conclusion could be drawn. Again, the cases follow the same two groupings as in the convergence history plots. The plots for the $\delta = 0.1, 0.25, 0.5$ and 0.75 cases are practically on top of each other, while the plots for C3P with explicit trailing edge region blanking and PEGASUS5 are very close to each other.

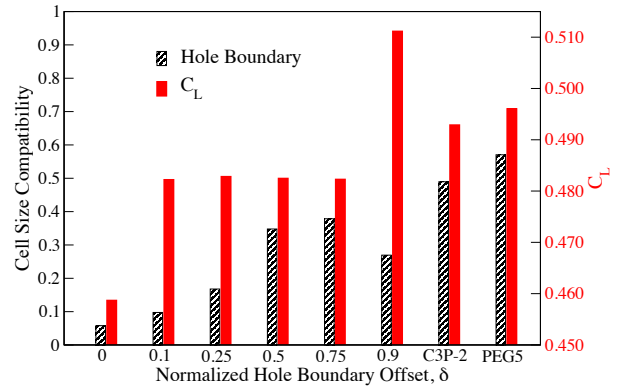
Figures 19a, b show the surface pressure coefficient difference between a representative case from group 1 (C3P with trailing edge region blanking), and a representative case from group 2 ($\delta = 0.5$). The plots illustrate the difference in the shock location on the upper surface of the wing. Figure 19c shows the C_p profile along a constant span cut on the outboard side of the wing. This clearly depicts that cases from group 2, by not resolving and treating the trailing edge region accurately, the circulation around the airfoil section and the shock location are slightly shifted from those from group 1. The lesson to take away from this test case is that if there are multiple grids resolving the same region in a high-gradient region near the geometry surface (e.g., leading and trailing edges of wings), the grid subsets that do not sufficiently resolve the flow should be blanked or treated as fringe points.

III.D. D8 Double Bubble Subsonic Aircraft

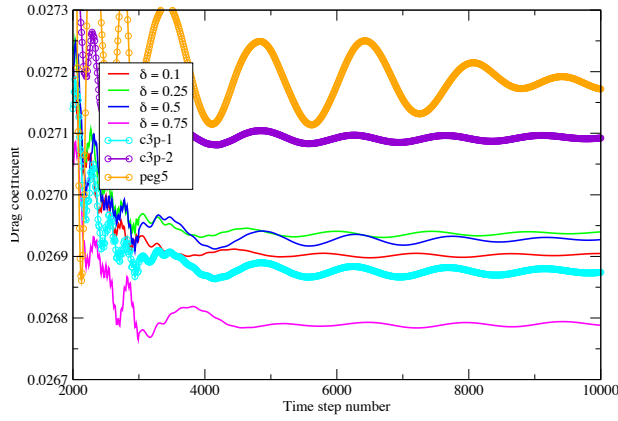
The D8 test case consists of a half-body subsonic aircraft with a fuselage, a wing, a vertical and a horizontal tail embedded in a wind tunnel with a support strut (Figure 20a). A small gap exists between the strut and the fuselage. This configuration contains both intersecting components (fuselage, wing, and tails),



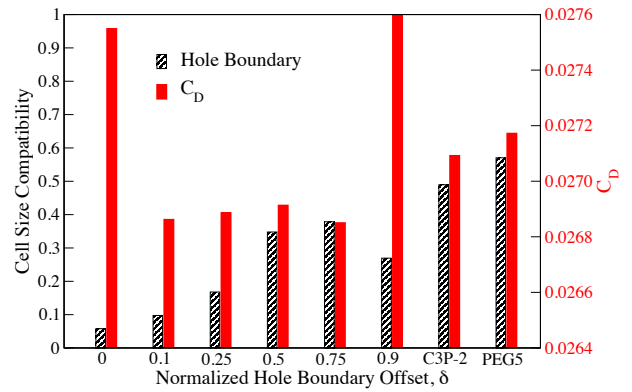
(a)



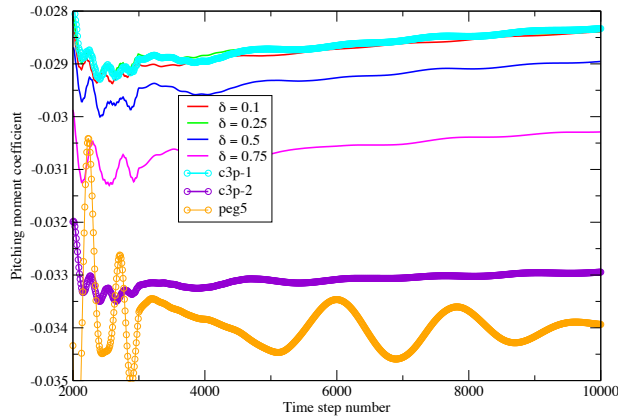
(b)



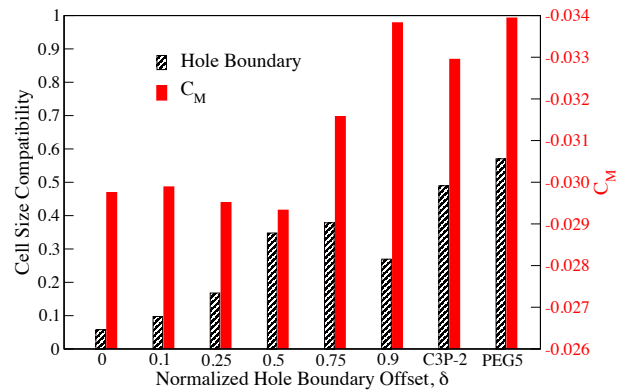
(c)



(d)



(e)



(f)

Figure 17. DPW4 test case aerodynamic loads convergence histories and hole boundary fringe points interpolation compatibility (c3p-1 = no trailing-edge region blanking, c3p-2 = with trailing-edge region blanking). (a) C_L convergence history. (b) Cell size compatibility and C_L . (c) C_D convergence history. (d) Cell size compatibility and C_D . (e) C_M convergence history. (f) Cell size compatibility and C_M .

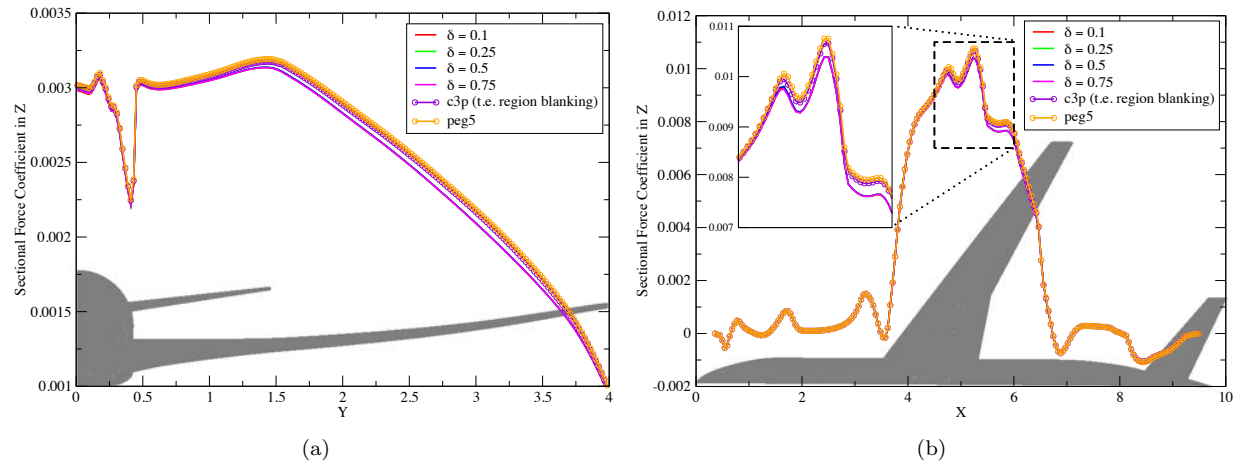


Figure 18. Sectional force coefficient in z for DPW4 test case. (a) Sectional force coefficient in z against y . (b) Sectional force coefficient in z against x (global view with zoom box marked by dashed lines).

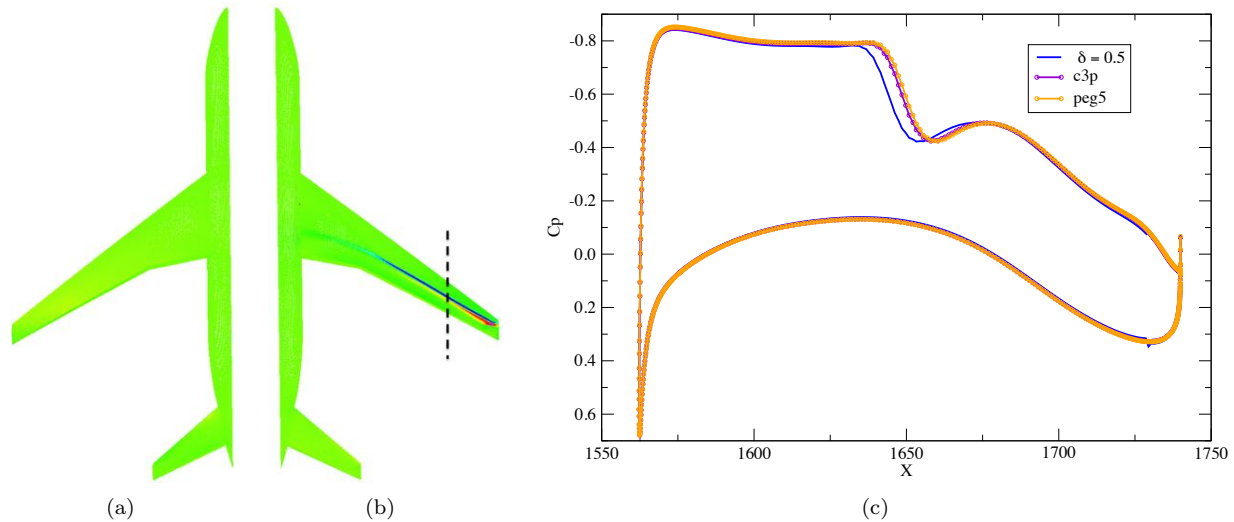


Figure 19. Difference in surface pressure coefficient C_p between C3P case with trailing edge blanking and $\delta = 0.5$ case for DPW4 (red indicates positive difference ≥ 0.1 , blue indicates negative difference ≤ -0.1 , green indicates approximately zero difference). (a) Bottom surface. (b) Top surface. Dotted line indicates constant span station cut for figure c. (c) C_p against streamwise distance x on constant span cut marked by dotted line in figure b.

and disjoint components (fuselage and strut). Curvilinear near-body grids are used to model the various components while several off-body stretched Cartesian box grids are employed to fill the interior of the wind tunnel. The entire grid system contains 35 grids and 83 million grid points. The solution is run at a freestream Mach number of 0.16, zero angle of attack, and a Reynolds number per inch of 79500.

Figure 20a shows various volume grid slices with automated hole boundaries created by C3P. Figure 20b shows the pressure coefficient contours for the converged solution on the surface geometry. The half-body solution has been reflected to show the solution on the full configuration so that the solution on the inboard side of the vertical tail is visible. The number of orphan points for different δ 's, C3P, and PEGASUS5 is shown in Table 4. All cases have either zero or a small number of orphan points ($< 0.03\%$) relative to the total number of grid points except for the $\delta = 0.9$ case which failed to converge due to the relatively large number of orphan points (0.3%). The 139 orphan points for the $\delta = 0$ case are due to the choice of the constant X-ray image plane spacing which is inadequate in resolving certain regions of the geometry. In practical applications of standard X-rays, $\delta = 0$ (minimum hole) is never used, and hence this situation never arises. The 12 orphan points for the PEGASUS5 case are located in the trailing edge region of the horizontal tail. Several highly twisted cells exist just off the surface of the razor sharp trailing edge which prevented the stencil search procedure from finding the appropriate donor cell. In practice, the OVERFLOW solver fills these orphan points with a reasonable solution that is averaged from their validly computed neighbors.

Table 4. Number of orphan points for various δ 's, C3P, and PEGASUS5 for D8 test case.

Hole boundary treatment	$\delta = 0.0$	$\delta = 0.1$	$\delta = 0.25$	$\delta = 0.5$	$\delta = 0.75$	$\delta = 0.9$	C3P	PEG5
No. of orphan points	139	1	1	46	21622	238313	0	12

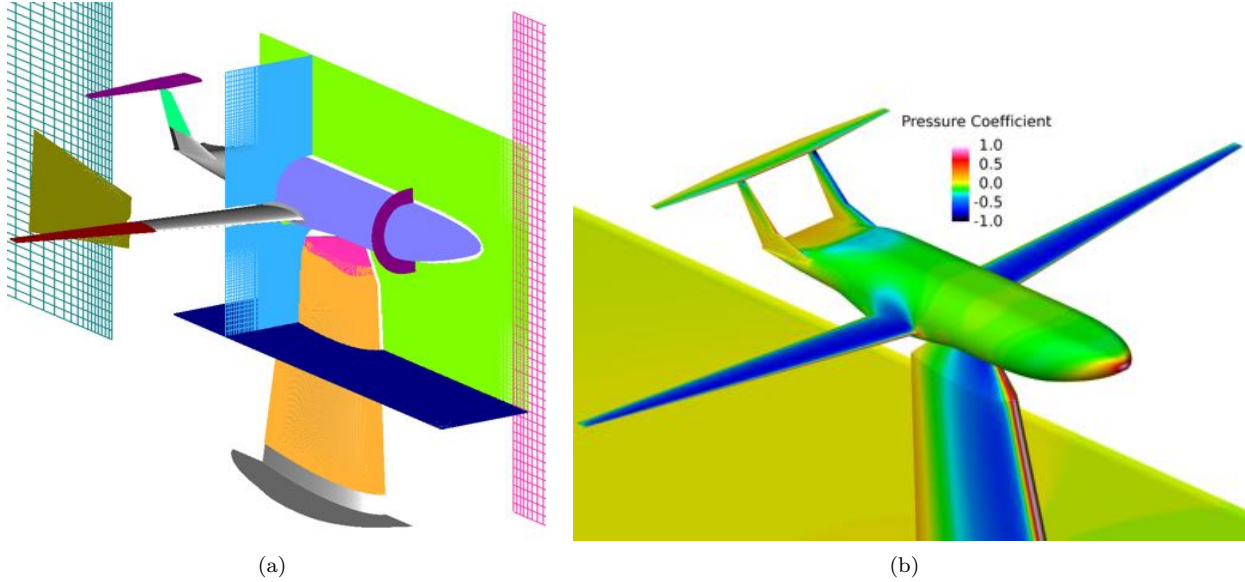


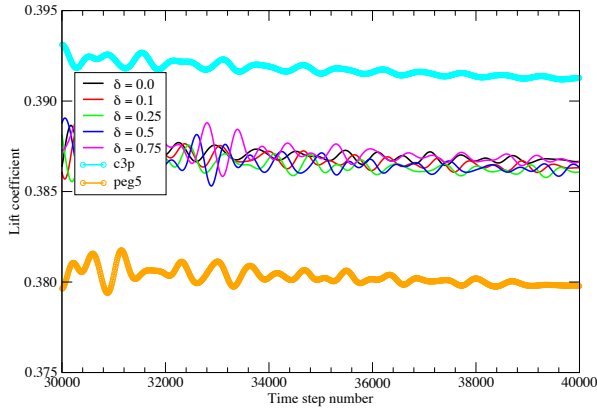
Figure 20. D8 test case. (a) Volume grid slices for near-body grids and off-body Cartesian grid showing hole boundary location created automatically by C3P. (b) Pressure coefficient contours on surface geometry of full-body obtained by reflecting half-body solution.

Figures 21a, 21c, 21e show the convergence histories of the lift (C_L), drag (C_D), and pitching moment (C_M) coefficients. Again, all cases appear to converge at about the same rate except the oscillations appear to damp out faster for the C3P case. For C_L and C_M , the C3P and PEGASUS5 cases are converging to different values than the δ cases. For C_D , the C3P case appears to be converging to a different value than the other cases but the difference is small (about 2 drag counts). The $\delta = 0.0$ case for C_D is the outlier here with a converged value of 0.0394 which lies outside the bounds of Figure 21c. The maximum deviation of the final converged values of all cases to the mean converged value is about 1.6% for C_L , 0.4% for C_D (excluding the $\delta = 0$ case), and 2.2% for C_M . These variations are of the same order of magnitude as those found in a study on the effects of various numerical methods, grid parameters, and turbulence models on the D8.¹⁴ For example, by varying the grid stretching ratio from 1.1 to 1.25, variations in C_L and C_D range from 1% to 12.5%. Usage of the Spalart-Allmaras (SA),¹⁸ Baldwin-Barth¹⁹ and SST²⁰ turbulence models show variations in C_L and C_D of 2.1% - 5.5% relative to the SA model.

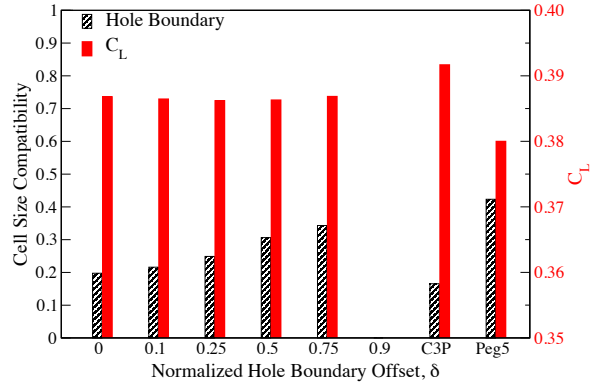
Figures 21b, 21d, 21f show the cell size compatibility between hole boundary fringe points and donor cells for various δ 's, C3P, and PEGASUS5. The PEGASUS5 case appears to have the best compatibility value in this case but the final converged loads for the different cases do not appear to correlate with the compatibility value.

Similar to the DPW4 test case, differences in the hole boundary and fringe point treatments on the wing leading and trailing edge regions lead to slight differences in the circulation and the surface pressure coefficient on the wing surface. Such small differences are illustrated in a constant span cut plot of C_p against x in Figure 22. The $\delta = 0.5$ curve falls between the C3P and PEGASUS5 curves, but all three are very close to each other.

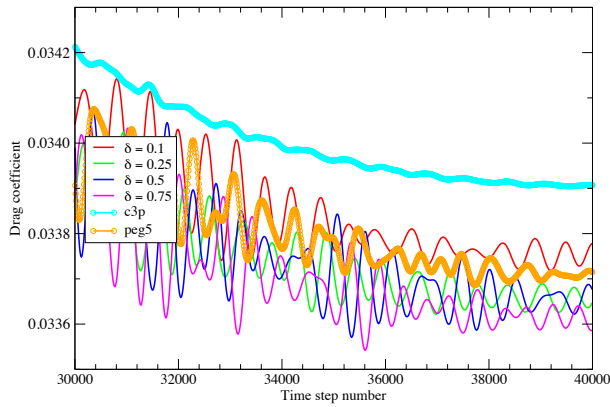
Figures 23 compares the different hole boundary treatments around the wing by showing grid slices with iblanks and fringe points (symbols) along a constant span cut. There are three grids in the vicinity of the wing: a body-fitted curvilinear O-grid around the wing (orange), a wake box grid that goes right up to the trailing edge of the wing (blue), and an off-body Cartesian box grid (green). In the leading edge region, the effective hole boundary from PEGASUS5 is further away from the surface compared to that from OVERFLOW-DCF ($\delta = 0.5$) and C3P. PEGASUS5 adds second-level fringe points to the entire Cartesian grid region in favor of the wing C-grid, creating a virtual hole in this box grid that is bigger than either the $\delta = 0.5$ or the C3P grids.



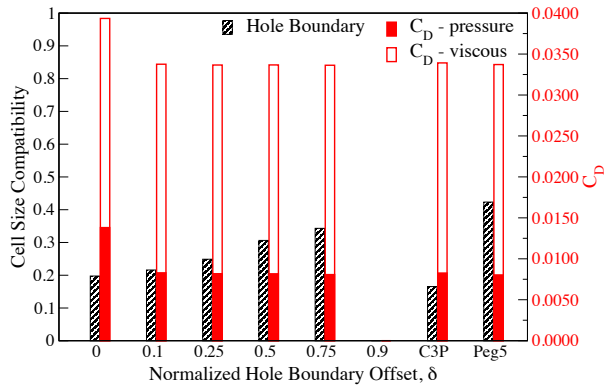
(a)



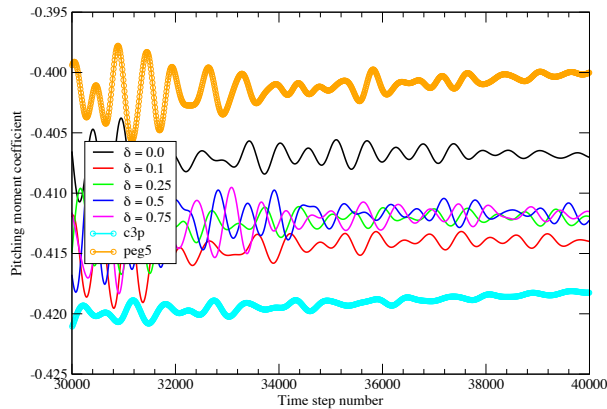
(b)



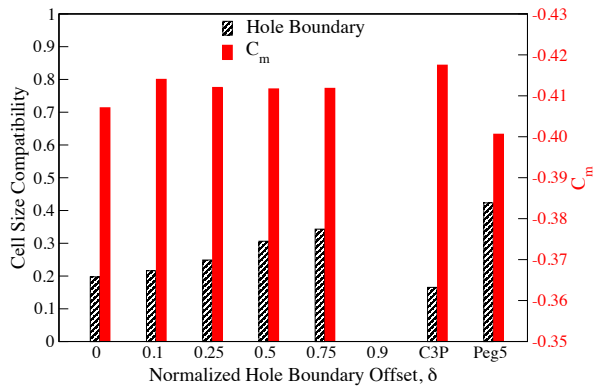
(c)



(d)



(e)



(f)

Figure 21. D8 test case aerodynamic loads convergence histories and hole boundary fringe points interpolation compatibility. (a) C_L convergence history. (b) Cell size compatibility and C_L . (c) C_D convergence history. (d) Cell size compatibility and C_D . (e) C_M convergence history. (f) Cell size compatibility and C_M .

In the trailing edge region, the manual OVERFLOW-DCF scheme allows both the near-body wing grid points and the wake box grid points to remain as field points except for the double fringe outer boundary of the wake box grid (Figure 23j). The off-body Cartesian grid has been blanked in this region. For C3P, both the wake box grid and the off-body Cartesian grid are blanked in the immediate vicinity of the trailing edge, leaving the near-body wing grid as the only grid used to resolve the flow immediately behind the trailing edge (Figure 23k). A little further downstream, the wake grid takes over in resolving the wake (Figure 23h). For PEGASUS5, most of the near-body wing grid points in the wake are treated as second level fringe points except for a few layers of points off the wing surface. The wake box grid points are left as field points in this region except for the double fringe outer boundary at the upstream end (Figure 23l). Of the three schemes, the hole boundary of the off-body Cartesian grid is closest to the wing surface for C3P.

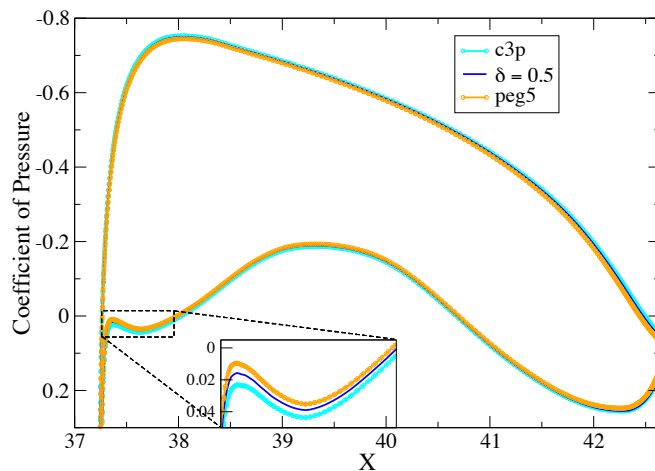


Figure 22. Surface pressure coefficient against streamwise distance x on a constant span cut on the wing for the D8 test case.

IV. Summary and Conclusions

The most expensive step in the automated X-rays approach¹² for determining hole boundaries in overset grids was the computation of an auxiliary wall-distance function. An efficient approximate method to compute this function is outlined in the current work. This involves determining an accurate wall-distance for cut-cells in a reference Cartesian grid, followed by a Fast Marching Method to fill the remaining grid points in the Cartesian mesh with an approximate wall distance. Wall distance for an arbitrary grid point is then easily obtained by interpolation from the reference Cartesian mesh. Results from test cases indicate that the new method, when implemented into the domain connectivity code C3P, is able to produce connectivity results in a comparable amount of wall-clock time to the original X-rays method in the OVERFLOW-DCF software for most of the test cases, despite additional computational work.

The sensitivity of converged values and convergence rates of aerodynamic loads to hole boundary locations is studied using four test cases. These include an isolated re-entry capsule in supersonic flow, a two-rockets configuration in supersonic flow, the AIAA 4th Drag Prediction Workshop CRM in transonic flow, and the D8 Double Bubble subsonic aircraft. For each test case, a sequence of runs are performed with different normalized hole boundary offset locations created by manual specification using the original X-rays approach, with the automatically-created hole boundary location from the improved X-rays scheme in C3P, and with the level-2 fringe treatment using a cell-difference parameter in PEGASUS5. At each hole boundary location, a measure of average cell volume compatibility is also evaluated between all the fringe points and their donor cells. Results from the runs indicate that the convergence rates of the aerodynamic loads appear to be about the same for all hole boundary locations, but the converged values vary slightly. For all test cases, variations in the converged loads between different hole boundary locations in the ‘best practice’ range ($0.25 \leq \delta \leq 0.75$), and for C3P and PEGASUS5 are small relative to effects from other sources such as grid stretching and turbulence models. Similarly, no obvious correlations are found between reasonable values of average cell size compatibility (≥ 0.2) and final converged loads for the test cases presented.

The observations above suggest that an averaged measure of cell attribute compatibility is too insensitive to distinguish differences between the effects of different hole boundary locations. A more relevant measure of compatibility between fringe points and their donor stencils may require considerations of both grid attributes and solution gradients. A significant fraction of the observed differences in aerodynamic loads is most likely originating from a relatively small region of the flow field. This is demonstrated by the DWP4 test case where differences in hole boundary locations and fringe points treatment occur in a small part of the domain with high gradients (e.g., leading and trailing edges of wings), which lead to small differences in the final converged aerodynamic loads. A recommended best practice from these results is that when multiple

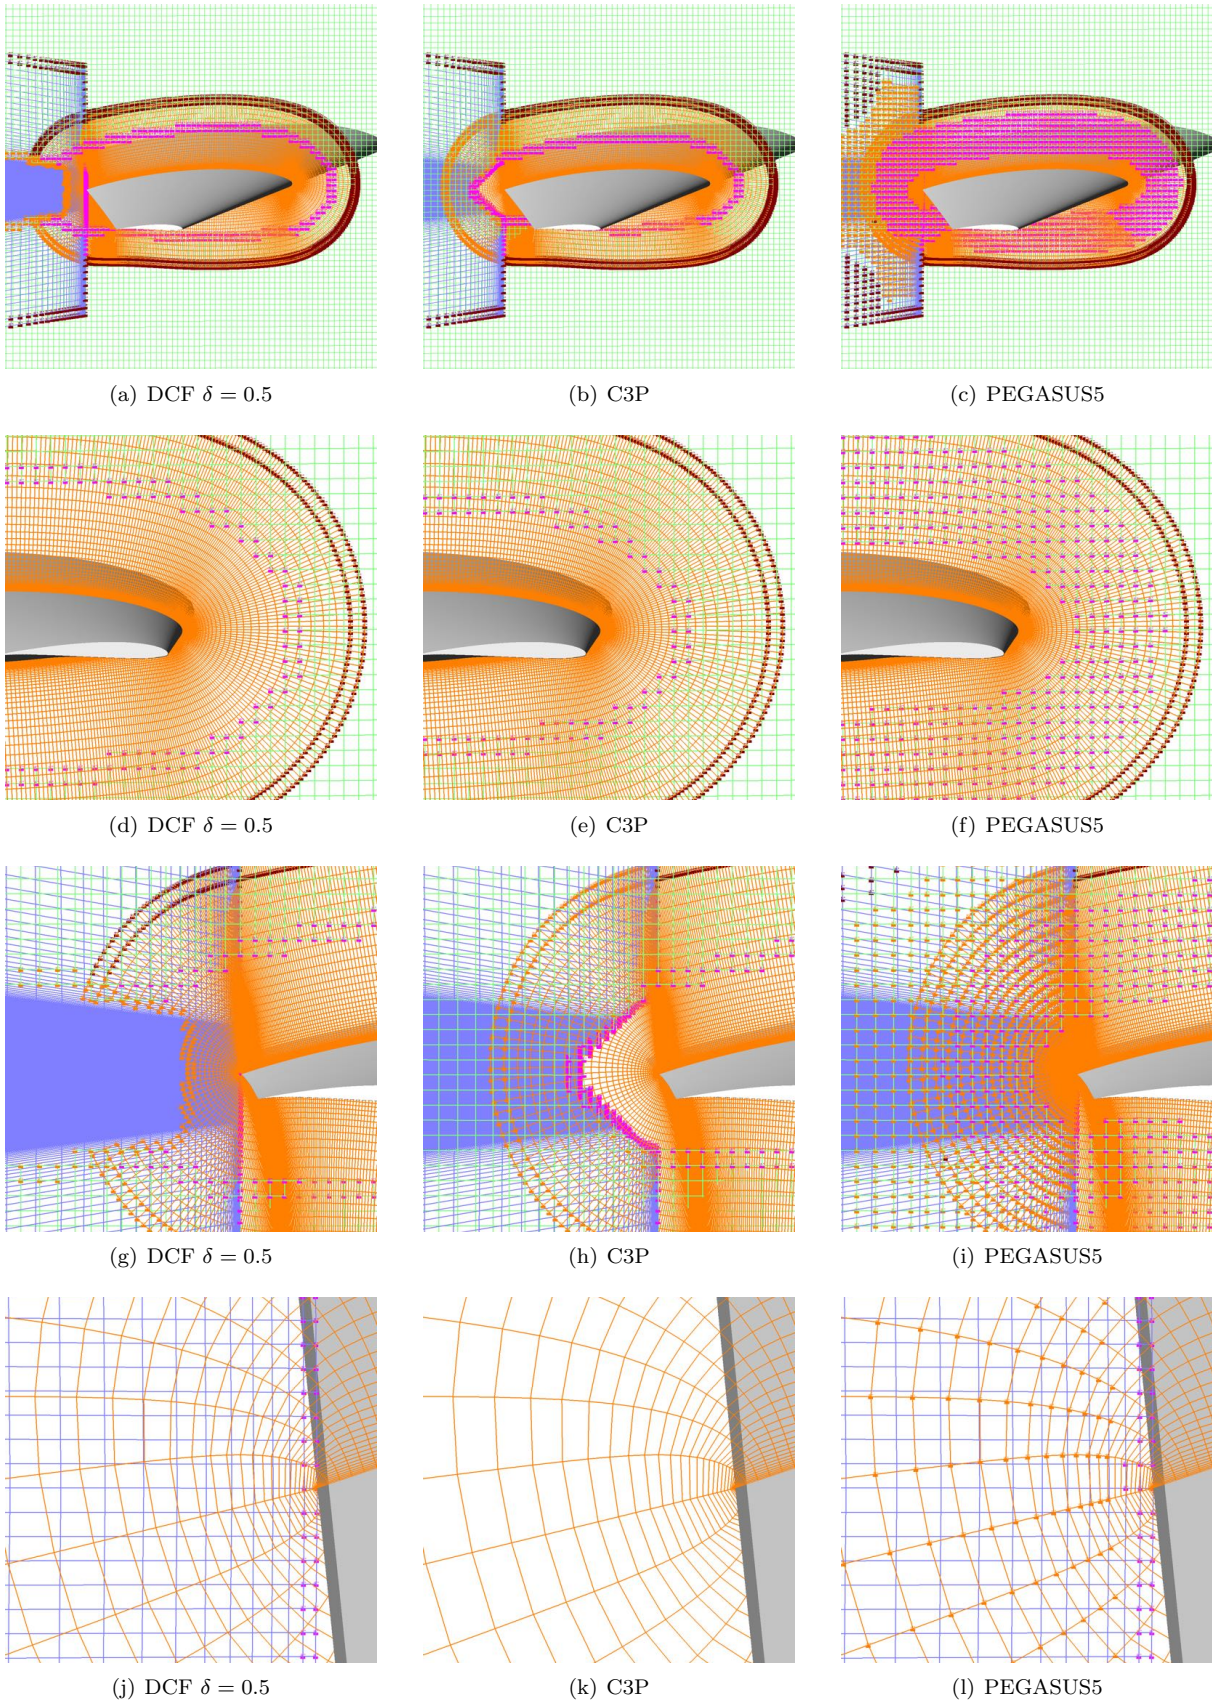


Figure 23. Hole boundaries on span cut through wing in D8 test case. Fringe points are marked by symbols. (a), (b), (c) Wing view. (d), (e), (f) Leading edge view. (g), (h), (i) Trailing edge overall view. (j), (k), (l) Trailing edge zoomed view.

overlapping grids do not all resolve the flow equally well, grid subsets from the coarser grids should not be left as field points to influence the flow but should be blanked or treated as second level fringe points. This is automatically performed in PEGASUS5, but currently manually specified in C3P. Further improvements in C3P could be implemented to consider both cell attribute compatibility and orphan points removal as objectives in the hole boundary iteration step.

The current work has not addressed relative motion problems where the original X-rays scheme has an efficiency advantage over other methods. As with the original X-rays map, the adaptive X-rays map¹² is almost as efficient in cutting the minimum hole during relative motion since the same coordinate transform is used in both cases for the Cartesian map look-up. Further analysis is needed to determine if a wall distance function can be easily updated from an initial state and a motion transform, and if similar efficiencies are attainable for orphan points iterations. These are issues that could be examined in future work.

Acknowledgements

The authors would like to thank Michael Aftosmis from NASA Ames Research Center for pointing out the merits of the Fast Marching Method for computing distances to the wall, and Dr. Jeffrey Housman from NASA Ames Research Center for insightful discussions on various aspects of hole cutting.

References

- ¹Gomez, R. J., Vicker, D., Rogers, S. E., Aftosmis, M. J., Chan, W. M., Meakin, R. L. and Murman, S., "STS-107 Investigation Ascent CFD Support," AIAA Paper 2004-2226, 2004.
- ²Ahmad, J. U., Pandya, S. A., Chan, W. M. and Chaderjian, N. M., "Navier-Stokes Simulation of Air-Conditioning Facility of a Large Modern Computer Room," FEDSM 2005-77225, Proceedings of the 2005 ASME Fluids Engineering Division Summer Meeting and Exhibition, Houston, Texas, 2005.
- ³Pandya, S., Onufer, J., Chan, W. and Klopfer, G., "Capsule Abort Recontact Simulation," AIAA Paper 2006-3324, 2006.
- ⁴Kiris, C. C., Kwak, D., Chan, W. M., Housman, J. A., "High-Fidelity Simulations of Unsteady Flow Through Turbopumps and Flowliners," *Computers & Fluids*, Vol. 37, pp. 536-546, 2008.
- ⁵Bhagwat, M., Dimanlig, A., Saberi H., Meadowcroft, E., Panda, B. and Strawn, R., "CFD/CSD Coupled Trim Solution of the Dual-Rotor CH-47 Helicopter Including Fuselage Modeling," Proceedings of the American Helicopters Society Aeromechanics Specialist's Conference, San Francisco, 2008.
- ⁶Kiris, C., Housman, J., Gusman, M., Chan, W. and Kwak, D., "Time-Accurate Computational Analysis of Ignition Overpressure in the Flame Trench," *Computational Fluid Dynamics Review*, Eds. Hafez, Oshima, Kwak, Publisher: World Scientific, 2010.
- ⁷Rogers, S. E., Suhs, N. E. and Dietz, W. E., "PEGASUS5 : An Automated Pre-Processor for Overset-Grid CFD," *AIAA J.*, Vol. 41, No. 6, pp. 1037-1045, Dec. 2003.
- ⁸Noack, R. W., "Suggar++: An Improved General Overset Grid Assembly Capability," AIAA Paper 2009-3992, 2009.
- ⁹Sitaraman, J., Floros, M., Wissink, A. and Potsdam, M., "Parallel Domain Connectivity Algorithm for Unsteady Flow Computations Using Overlapping and Adaptive Grids," *J. Comp. Phys.*, Vol. 229, pp. 4703-4723, March 2010.
- ¹⁰Henshaw, W. D., "Ogen: An Overlapping Grid Generator for Overture," Research Report UCRL-MA-132237, Lawrence Livermore National Laboratory, 1998.
- ¹¹Meakin, R. L., "Object X-rays for Cutting Holes in Composite Overset Structured Grids," AIAA Paper 2001-2537, 2001.
- ¹²Chan, W. M., Kim, N. and Pandya, S. A., "Advances in Domain Connectivity for Overset Grids Using the X-rays Approach," Paper ICCFD7-1201, 7th International Conference on Computational Fluid Dynamics, Big Island, Hawaii (http://www.iccfd.org/iccfd7/assets/pdf/papers/ICCFD7-1201_paper.pdf), 2012.
- ¹³Vassberg, J. C., Tinoco, E. N., Mani, M., Rider, B., Zickuhr, T., Levy, D. W., Brodersen, O. P., Eisfeld, B., Crippa, S., Wahls, R. A., Morrison, J. H. Mavriplis, D. J., Murayama, M., "Summary of the Fourth AIAA CFD Drag Prediction Workshop," AIAA Paper 2010-4547, 2010.
- ¹⁴Pandya, S. A., "External Aerodynamics Simulations for the MIT D8 "Double-Bubble" Aircraft Design," Paper ICCFD7-4304, 7th International Conference on Computational Fluid Dynamics, Big Island, Hawaii (http://www.iccfd.org/iccfd7/assets/pdf/papers/ICCFD7-4304_paper.pdf), 2012.
- ¹⁵Hassouna, M. S., Farag, A. A., "Multistencils Fast Marching Methods: A Highly Accurate Solution to the Eikonal Equation on Cartesian Domains," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 29, No. 9, pp. 1-12, 2007.
- ¹⁶Slotnick, J. P., Hannon, J. A. and Chaffin, M., "Overview of the First AIAA CFD High Lift Prediction Workshop," AIAA Paper 2011-0862, 2011.
- ¹⁷Nichols, R. H., Tramel, R. W. and Buning, P. G., "Solver and Turbulence Model Upgrades to OVERFLOW 2 for Unsteady and High-Speed Applications," AIAA Paper 2006-2824, 2006.
- ¹⁸Spalart, P. R., Allmaras, S. R., "A One-Equation Turbulence Model for Aerodynamic Flows," AIAA Paper 1992-0439, 1992.
- ¹⁹Baldwin, B. S. and Barth, T. J., "A One-Equation Turbulence Transport Model for High Reynolds Number Wall-Bounded Flows," NASA TM 102847, 1990.
- ²⁰Menter, F. R., "Two-Equation Eddy-Viscosity Turbulence Models for Engineering Applications," *AIAA J.*, Vol. 32, No. 8, pp. 1598-1605, 1994.